“World Targets in Megadeaths” and Stabilizing the U.S.–China Rivalry: From RAND’s Nuclear Escalation Doctrine to Managed Great-Power Competition

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

November 22, 2025

In the early Cold War, RAND analysts such as Herman Kahn helped teach Washington how to think about thermonuclear war in units of “megadeaths.” Cities became entries on targeting lists; millions of civilian deaths were converted into curves, ladders, and “tragic but distinguishable postwar states.” This mindset was seared into popular memory by Stanley Kubrick’s Dr. Strangelove, where the fictional “World Targets in Megadeaths” binder distilled the icy logic of deterrence and escalation. Seventy years later, RAND has produced a very different kind of document on the central rivalry of the age: a report on stabilizing U.S.–China competition that emphasizes mutual restraint, crisis management, and acceptance of long-term coexistence. It quietly abandons the language of decisive victory and calls instead for limiting damage in three high-risk domains: Taiwan, the South China Sea, and science and technology. The report’s subsequent withdrawal for “further review” highlights just how politically fragile this shift has become. This paper reads the two RAND eras together. It traces the evolution from megadeath arithmetic to interdependence management, asks why a stabilization blueprint is now awkward to leave on the record, and offers a logostic interpretation of the attraction toward, and resistance to, a τ that is “not war.”

Keywords: RAND Corporation, Herman Kahn, megadeaths, Dr. Strangelove, nuclear deterrence, U.S.–China rivalry, Taiwan, South China Sea, science and technology competition, semiconductors, escalation management, strategic culture, great-power politics, interdependence, logostics, q_US(t), q_PRC(t), τ_not_war, ethics of deterrence.

A collaboration with GPT-5.1-Thinking-Extended. 35 pages. CC4.0.

Outline

1. Introduction – Two RANDs, One Empire
1.1 Setting the puzzle: from megadeath curves to stabilization strategies
1.2 Why the withdrawal of a U.S.–China report matters
1.3 Method and perspective of this paper

2. RAND, Herman Kahn, and the Megadeath Mindset
2.1 The birth of nuclear strategy and planning in megadeaths
2.2 Escalation ladders, "distinguishable postwar states," and rational catastrophe
2.3 Dr. Strangelove as cultural x-ray of Cold War strategic thinking
2.4 How the megadeath logic seeped into institutions and public imagination

3. The "Stabilizing the U.S.–China Rivalry" Report
3.1 Origins, structure, and stated purpose of the report
3.2 Taiwan: deterrence without independence, incentives for gradualism
3.3 South China Sea: allies, red lines, and avoiding third-party escalation
3.4 Science and technology: infrastructure, semiconductors, and AI as new critical targets
3.5 The quiet withdrawal: "further review" and its possible meanings

4. From World Targets in Megadeaths to Interdependent Vulnerabilities
4.1 Replacing city lists with fabs, cables, ports, and platforms
4.2 Economic entanglement and mutual hostage situations
4.3 Information, finance, and supply chains as non-kinetic escalation ladders
4.4 Why contemporary war planning cannot be separated from global systems survival

5. Strategic Culture and Political Red Lines in Washington
5.1 Victory narratives, moral exceptionalism, and domestic electoral pressures
5.2 Hawks, doves, and the suspicion of "legitimizing" the Chinese Communist Party
5.3 Taiwan as symbol, ally, and tripwire for credibility politics
5.4 Why candid stabilization language is harder to publish than megadeath tables ever were

6. A Logostic Reading: q_US(t), q_PRC(t), and τ_not_war
6.1 Megadeath doctrine as one historical attractor for q_US(t)
6.2 Managed rivalry and interdependence as a competing attractor
6.3 Curvature, humility, and the acceptance of finite power under τ_not_war
6.4 How the withdrawal reflects friction between institutional q and any stabilizing τ

7. After Megadeaths: Ethics, Policy, and the Possibility of Repentance

7.1 Minimal ethical conditions for honest great-power strategy today

7.2 Rethinking alliances, extended deterrence, and Taiwan under a non-victory paradigm

7.3 Guardrails for science and technology competition that stop short of megadeath by other means

7.4 Roles for citizens, churches, and dissident networks in resisting escalation culture

8. Conclusion – What RAND's Two Eras Reveal

8.1 Summary of the argument and its limits

8.2 Open questions for strategists, ethicists, and theologians

8.3 Final reflection: living under τ_not_war in a world still haunted by megadeath thinking

9. References

1. Introduction – Two RANDs, One Empire

The RAND Corporation has been in the room for almost every major American debate about war, peace, and "acceptable" catastrophe since the dawn of the nuclear age. In the 1950s and 1960s, RAND analysts helped invent a new language for thinking about thermonuclear war: escalation ladders, counterforce, second-strike capability, and, notoriously, "megadeaths" as a unit of planning. The idea that millions of civilian deaths could be plotted on curves and weighed against each other entered the bloodstream of policy and popular culture alike. Stanley Kubrick's Dr. Strangelove parodied this world with its infamous "World Targets in Megadeaths" binder, but the joke worked only because it was so recognizably close to reality.

Decades later, RAND has produced a very different kind of document about the central great-power rivalry of our own time: a report on stabilizing the U.S.–China relationship. Instead of thinking in terms of decisive victory and acceptable levels of destruction, it speaks the language of mutual restraint, crisis management, and damage limitation across three interlocking domains: Taiwan, the South China Sea, and science and technology. Rather than treating "the enemy" as an object to be broken, it treats a rival superpower as a permanent feature of the landscape that must be lived with, not destroyed.

Then, without public explanation, this stabilization report is withdrawn from RAND's own website for "further review." The PDF continues to circulate, but the institutional door is quietly half-closed. The juxtaposition is striking: in one era, the most chilling megadeath charts could sit on a public bookshelf; in another, a pragmatic blueprint for coexistence becomes awkward to leave on the front page. The existence, content, and retreat of this document form the starting point for this paper.

1.1 Setting the puzzle: from megadeath curves to stabilization strategies

The puzzle can be stated simply. Early RAND was willing to think and write in brutally explicit terms about nuclear Armageddon. Its analysts described thermonuclear exchanges, fallout patterns, and "tragic but distinguishable postwar states" with a clinical calm that horrified many outside the community. Yet those texts were not memory-holed. They were published, debated, and even mocked in film, but they remained part of the official, citable record of American strategic thought.

By contrast, the contemporary report on U.S.–China stabilization does not call for surrender, nor does it celebrate the Chinese Communist Party. It does something more modest and, in some ways, more unsettling: it accepts that there is no clean path to victory that does not carry intolerable systemic risks. It accepts the long-term survival of a peer adversary and explores concrete ways to reduce the odds of catastrophe in key theaters. It nudges the reader toward a

world where “winning” looks less like defeating an enemy and more like never letting certain escalations happen at all.

Why should this second posture be the one that cannot stand unchallenged on RAND’s website? How can a political system that once tolerated megadeath arithmetic now struggle to leave a stabilization blueprint in plain sight? The transition from explicit war-gaming in units of millions of dead to carefully hedged coexistence strategies is the core tension this paper explores.

1.2 Why the withdrawal of a U.S.–China report matters

On one level, the withdrawal of a single think-tank report might look trivial. RAND routinely produces studies for different government sponsors; drafts are revised, delayed, or shelved for all sorts of technical or bureaucratic reasons. Yet the public release and quiet retreat of a document explicitly dedicated to stabilizing the U.S.–China rivalry is symbolically charged in at least three ways.

First, it illuminates the current political red lines in Washington. Explicitly accepting the long-term legitimacy of the Chinese regime, acknowledging that there is no plausible path to “absolute victory,” and recommending subtle restraint on the part of allies and partners all cut against popular narratives of toughness and moral clarity. Leaving such recommendations in a prominent place risks making visible an internal realism that much of the political class prefers to downplay.

Second, the withdrawal matters because the report grapples with domains in which escalation would be civilizationally costly even if no nuclear weapons were used. Taiwan’s semiconductor fabs, the sea lanes and cables of the Western Pacific, and the intertwined digital and physical infrastructure of science and technology are new “world targets” whose destruction would cripple the global system. Publicly admitting that both sides must exercise restraint around these vulnerabilities exposes how fragile the shared platform of late modern civilization has become.

Third, the episode offers a window into institutional self-consciousness. If the megadeath era was marked by a kind of icy candor, the stabilization era seems marked by a nervous ambivalence: the need to think about stabilizing rivalry coexists with the fear of being accused of appeasement or defeatism. The disappearance of the report from RAND’s front page is not just a bureaucratic footnote; it is a small but telling sign of how constrained serious discussion of non-victory outcomes has become.

1.3 Method and perspective of this paper

This paper approaches the “two RANDs” contrast in three intertwined ways. First, it offers a historical and conceptual comparison of strategic cultures. It reads early nuclear-era texts and

the contemporary stabilization report as snapshots of how American elites imagine catastrophe, victory, and acceptable risk. The focus is less on technical details and more on the moral geometry implicit in phrases like “tragic but distinguishable postwar states” or “sustained coexistence with a rival superpower.”

Second, the paper treats the withdrawn U.S.–China report as a diagnostic artifact. Its content, its emphasis on mutual restraint, and its subsequent retreat from public prominence are interpreted as symptoms of a deeper tension in U.S. grand strategy: a reluctance to say openly that the empire may need to live indefinitely with limits, peers, and permanent vulnerability. In this sense, the report’s fate is as important as its recommendations.

Third, the analysis uses a logostic lens. It interprets U.S. strategic postures as time-dependent states q_US(t) and Chinese postures as q_PRC(t), both moving within and against a higher-order τ_not_war: a truth manifold in which full-scale great-power war, whether nuclear or systemic-technological, is recognized as incompatible with continued human flourishing. The comparison between megadeath doctrine and stabilization doctrine thus becomes more than a technical shift. It becomes a question about whether major powers can bend their institutional trajectories toward a τ that refuses to treat the destruction of millions, or of the world’s infrastructure, as an acceptable cost of “credibility.”

This perspective does not claim omniscience and does not pretend to be neutral. It openly favors a world in which τ_not_war is taken seriously as a binding constraint. It asks what it means that earlier generations could openly publish their megadeath tables, while a document trying to avoid twenty-first-century versions of those tables must be half-hidden.

2. RAND, Herman Kahn, and the Megadeath Mindset

Before there could be a report on “stabilizing” a great-power rivalry, there had to be a way to think about destroying the world without flinching. In the first decades of the nuclear age, RAND was one of the key laboratories where that mental apparatus was built. The task, as framed by its sponsors, was paradoxical: to prevent nuclear war by making it thinkable. That meant quantifying mass death, turning cities into datapoints, and drawing maps of catastrophe that could be read in a conference room.

Herman Kahn, perhaps more than any other single figure, embodies this early RAND ethos. Trained as a physicist and steeped in game theory, he approached thermonuclear war with a combination of technical sophistication and theatrical confidence. He spoke for a wider community of strategists who believed that rational analysis of “the unthinkable” was morally necessary: only by calculating potential losses, sequencing escalatory steps, and imagining

“survivable” nuclear exchanges could policymakers deter adversaries and avoid accidental annihilation.

The result was a new mental universe. In it, the death of millions could be represented as a point on a curve. “Acceptable” damage could be discussed not as a moral verdict but as a planning parameter. For a time, this felt to many insiders like realism—and to many outsiders like madness.

2.1 The birth of nuclear strategy and planning in megadeaths

Nuclear weapons arrived at the end of the Second World War with a shock that outstripped existing vocabulary. Early discourse oscillated between apocalyptic dread and vague hope that such destructive power would make war obsolete. RAND’s early nuclear strategists rejected both extremes. They insisted that nuclear weapons were not magic but instruments whose use, or non-use, could be disciplined by rational analysis.

To do that analysis, they needed units. Kilotons and megatons measured explosive yield; blast radii and fallout patterns measured physical effects. But policymakers cared about deaths, cities, and “national survival.” The bridge between physics and politics became the “megadeath”: one million human deaths treated as a single quantitative unit.

Thinking in megadeaths was not simply a stylistic choice. It allowed strategists to sketch trade-offs: Is it better to build more hardened silos or invest in missile defenses? How many millions dead on each side count as an intolerable loss? What combinations of first strikes, second strikes, and civil defense might produce a “tragic but not terminal” outcome? These questions were horrifying, but within the internal logic of deterrence they were presented as necessary. Without some willingness to measure, the argument went, there could be no way to signal to an adversary what was off-limits or to design forces that would make aggression look irrational.

2.2 Escalation ladders, “distinguishable postwar states,” and rational catastrophe

Herman Kahn’s escalation ladders formalized a certain style of thinking about crisis. Instead of imagining a single leap from peace to Armageddon, he broke conflict into discrete rungs: from minor diplomatic incidents to limited conventional skirmishes, to demonstration nuclear shots, to full-scale counterforce and countervalue exchanges. The purpose was to show that leaders might climb partway up the ladder, pause, and climb back down, rather than tumbling automatically into all-out war.

This same structural imagination gave rise to the language of “distinguishable postwar states.” Not every nuclear war was assumed to be the end of human civilization. Some scenarios produced “only” a few million deaths and left major power centers intact; others produced tens or hundreds of millions of deaths but left enough surviving infrastructure for a mangled version

of national life. These were not described as desirable outcomes. They were described as tragic, yet still meaningfully different from absolute annihilation.

The intended ethical justification was that only by thinking through such gradations could one design deterrent postures that made the worst outcomes unlikely. Critics justly noted that the exercise risked normalizing the unthinkable. Once “20 million dead” appears on a chart next to “50 million dead,” it invites a subtly perverse question: which is “better,” and under what conditions would a rational decision-maker accept one rather than the other? Nuclear strategy thus became the art of rational catastrophe: designing arsenals and doctrines that promised horrific costs in order to avert even worse ones.

2.3 Dr. Strangelove as cultural x-ray of Cold War strategic thinking

Stanley Kubrick’s Dr. Strangelove did not emerge from nowhere. It is, among other things, a darkly accurate portrait of the world that megadeath planners had built for themselves at RAND and similar institutions. The fictional war room, with its massive maps and confident briefings about “acceptable” casualty figures, exaggerates for satire but does not fundamentally distort. The famous line about “no more than ten to twenty million killed, tops” captures the obscene poise with which some strategists spoke about human losses.

The gag of the “World Targets in Megadeaths” binder works because it compresses the entire worldview into one absurd prop. Here is the planet, carved into targets, each associated with a number expressing expected deaths in millions. The visual joke is brutal, but it is also an x-ray: it shows how easily the language of megadeaths allows real places and people to vanish into abstract aggregates.

Importantly, the film does not depict its generals and experts as comic-book villains. They are absurd, compromised, sometimes monstrous, but they are also true believers in a certain concept of rationality. That is what makes the satire so sharp. It reveals a system in which intelligent, patriotic professionals have convinced themselves that they must be ready to incinerate cities in order to save them.

2.4 How the megadeath logic seeped into institutions and public imagination

Over time, the megadeath mindset migrated from specialized reports into the wider institutional and cultural fabric. In the Pentagon and allied ministries, planners used casualty tables, damage expectancy models, and targeting lists that treated entire regions as variables in optimization problems. In academia, game theory and systems analysis trained generations of analysts to think in terms of payoffs, equilibria, and utility maximization at scales that included millions of lives.

Publicly, the specifics of war plans remained classified, but the outline of the logic was well known. Civil defense campaigns, debates over missile gaps and bomber gaps, and arguments about "escalation dominance" all carried fragments of the megadeath calculus into mainstream politics. Citizens were told that deterrence worked because adversaries believed that the United States was both capable of mass destruction and rational enough to use that capability under certain conditions.

Popular culture absorbed and refracted these ideas. Films, novels, and television shows filled with mushroom clouds and last-minute hotline calls made nuclear strategy a familiar backdrop. The very phrase "mutual assured destruction" became a dark punchline, something schoolchildren could recite without ever having read a technical report. The underlying pattern was clear: apocalyptic risk was domesticated into a permanent feature of late modern life.

By the time the Cold War ended, the megadeath logic had done its work. It had helped hold the nuclear peace, but it had also trained institutions to think of existential threats as variables to be managed rather than boundaries never to be approached. That habit of mind did not disappear with the Soviet Union. It lingered in the background of strategic thinking, ready to be repurposed for new domains: cyber warfare, space, biotechnology, and now the deeply entangled relationship between the United States and China. The contemporary stabilization report emerges from this long shadow, trying to take seriously the costs of catastrophe without reverting to the language of "distinguishable postwar states" in a world where the targets are no longer only cities but the global systems that keep everything running.

3. The "Stabilizing the U.S.–China Rivalry" Report

The withdrawn RAND report appears at first glance as a conventional strategic study: authoritative authors, a familiar RAND format, a sober assessment of risks, and a menu of policy recommendations. It was written by Michael J. Mazarr, Amanda Kerrigan, and Benjamin Lenain, released in mid-October 2025 as a research report under RAND's National Security Research Division, and given a standard document number and DOI. The subject matter, however, is anything but routine. The report takes as its starting point the claim that the U.S.–China rivalry is structurally entrenched and globally consequential, and that the central task is not to "win" in any absolute sense but to keep this rivalry from spiraling into military conflict, economic warfare, or systemic fragmentation.

Structurally, the authors move from general principles to issue-specific applications. They first survey historical experience with stabilizing rivalries, distilling a set of broad conditions under which intense competition can be managed rather than allowed to explode. They then analyze Chinese strategic intent using recent studies of Chinese-language documents, asking how

Beijing itself frames the prospects for peaceful coexistence. Finally, they convene pairs of U.S. and Chinese scholars to explore what stabilization might look like in three concrete domains: Taiwan, the South China Sea, and competition in science and technology. The result is a blueprint that is less about dominance and more about building a limited, fragile modus vivendi within very specific danger zones.

3.1 Origins, structure, and stated purpose of the report

The report emerges from a particular moment in U.S. strategic debate. After years of "great-power competition" rhetoric, export controls, technology bans, and military signaling in the Western Pacific, senior officials and analysts increasingly acknowledge two hard facts: China is not going away, and the structural incentives for rivalry are not going to vanish. Against this background, RAND's project asks a pointed question: if the rivalry is durable, can it at least be stabilized?

According to the publicly available summaries and monitoring notes, the study proceeds in three main steps. First, it examines generic problems of stabilizing strategic rivalries—drawing on past examples to identify principles such as accepting a limited modus vivendi, mutually recognizing each other's political legitimacy, and building shared rules and norms in specific issue domains. Second, it assesses Chinese strategic intent, synthesizing the existing literature on Beijing's doctrinal writings to gauge whether China's leadership sees coexistence as possible or seeks fundamental revision of the international order. Third, it organizes work by paired U.S. and Chinese scholars, each pair tasked with exploring concrete options for stabilization in one of the three focus areas.

The stated purpose is modest but consequential. The report does not propose a grand bargain, a new alliance architecture, or an end to competition. Instead, it argues that both sides have a shared interest in preventing "worst-case" trajectories: outright war over Taiwan, a violent incident in the South China Sea that escalates beyond control, or technological brinkmanship that targets the other's critical infrastructure. Stabilization, in this sense, means reducing the probability of catastrophic outcomes while accepting that the rivalry will continue.

3.2 Taiwan: deterrence without independence, incentives for gradualism

Taiwan is the report's most sensitive and controversial domain. Here, the authors argue that stabilization requires a delicate balance: convincing Beijing that any use of force would be costly and uncertain, while avoiding symbolic or legal moves that Beijing would interpret as U.S.-backed Taiwanese independence. They explicitly resist the language of "absolute victory," instead framing U.S. objectives as deterring war and preserving Taiwan's security and democratic system within a long-term context of rivalry management.

A key line in the summary is that "stabilizing the Taiwan issue should focus on creating the maximum incentive for Beijing to pursue gradual approaches toward unification." This sentence has drawn particular political fire, because it openly acknowledges unification as a possible end state and proposes to shape Chinese choices toward slow, negotiated pathways rather than forceful ones. Underneath that phrase lies a bundle of recommendations:

- Reinforcing deterrence by maintaining credible U.S. and allied defense capabilities in the region, including continued arms sales and training for Taiwan.
- Coupling this with clear political reassurances that the United States does not support de jure independence, in line with its longstanding "One China" policy.
- Encouraging predictable high-level dialogues across all three sides (Washington, Beijing, Taipei) to reduce the risk of miscalculation.
- Expanding Track 2 diplomacy—unofficial contacts among former officials, scholars, business leaders, and retired officers—to keep channels open even when formal relations are strained.

Some versions of the recommendations, as paraphrased in secondary discussions, hint at a bolder step: Washington quietly discouraging provocative actions by Taiwan's leadership that might force Beijing's hand, while simultaneously maintaining robust deterrence against any Chinese use of force. In other words, Taiwan is treated neither as a pawn nor as a fully autonomous free agent, but as a central variable in a three-sided stability equation. For critics who want an unambiguous commitment to defend Taiwan's political choices regardless of consequence, this structured restraint looks worryingly like accommodation.

3.3 South China Sea: allies, red lines, and avoiding third-party escalation

In the South China Sea, the report's concern is less about a deliberate major war decision and more about accidental or third-party escalation. The area is crowded with coast guards, fishing fleets, maritime militias, and naval vessels from multiple countries. Any collision, ramming incident, or warning shot can become fodder for nationalist outrage and alliance tests.

The authors propose combining deterrence of military escalation with intensified diplomacy designed to make peaceful dispute management the default expectation. Summaries emphasize several strands:

- Washington and Manila should specify, in more detail than current public statements, what kinds of attacks or incidents would trigger U.S.–Philippine Mutual Defense Treaty obligations. This is meant to reassure the Philippines while also reducing the risk that small-scale confrontations automatically activate great-power commitments.

- Beijing, for its part, is encouraged to clarify its own red lines and escalation thresholds, so that U.S. and allied planners can better understand which actions would be read as intolerable. The report suggests that the United States may sometimes need to press its own partners not to test those thresholds recklessly.
- The study favors codes of conduct and agreed protocols for air and maritime encounters, building on and extending existing arrangements, to prevent dangerous maneuvers from becoming incidents.
- It suggests reducing the publicity given to certain U.S. naval and air operations—such as some intelligence, surveillance, and reconnaissance flights or freedom-of-navigation operations—without necessarily reducing their frequency. The goal is to maintain operational presence while depriving domestic hawks on all sides of spectacle.

Here the core idea is that third parties, emboldened by alliance guarantees or domestic political needs, can drag large powers into crises neither side really wants. Stabilization, in this domain, is as much about managing friends as managing adversaries.

3.4 Science and technology: infrastructure, semiconductors, and AI as new critical targets

The report's third focus area recognizes that science and technology competition has become a central theater of the U.S.–China rivalry, with semiconductors, data centers, undersea cables, and AI research facilities functioning as new "critical targets." In a globalized, digitized system, attacks on these nodes—even without any kinetic exchange—can produce cascades of economic and societal damage comparable to traditional military strikes.

The authors argue that both sides have an interest in drawing some boundaries around this domain. The summaries highlight recommendations such as:

- Mutual commitments not to target each other's core semiconductor fabrication plants, electrical grids, or other critical technological infrastructure, at least outside declared wartime.
- Reciprocal restraint pledges in especially dangerous areas, such as gain-of-function biological research, AI-enabled offensive cyber capabilities, and tools explicitly built to compromise the other side's command-and-control systems.
- Reviving U.S.–China dialogues on emerging technologies, including AI, to rebuild at least minimal trust and shared situational awareness.
- Exploring limited cooperation in non-sensitive domains: "AI for good" projects, climate and clean energy technologies, and some forms of basic science where mutual benefit is obvious and security concerns are manageable.

- In the longer term, beginning joint conversations on catastrophic AI-loss-of-control scenarios and on norms to prevent uncontrolled technological arms races from destabilizing the rivalry itself.

In this framing, stabilization does not mean halting competition. Export controls, investment screening, and rival industrial policies all continue. What changes is that both sides acknowledge a shared interest in not turning the world's core technological and informational infrastructure into a battlefield, because the resulting damage would not stay confined to the adversary.

3.5 The quiet withdrawal: "further review" and its possible meanings

The most dramatic aspect of the report's life has not been its content but its disappearance from RAND's public catalog. The PDF remains accessible at its original URL, and external sites continue to host the full text or detailed summaries. Yet the official RAND landing page now consists of a brief description of the topic, a statement that "this document has been withdrawn for further review," and a notice that availability is "withdrawn." No public explanation, erratum, or timeline for resolution accompanies this change.

Several interpretations are possible. A benign reading would point to internal quality-control processes. Perhaps a sponsor raised technical concerns, classified information issues, or objections about how certain sources were characterized. RAND's note that its publications do not necessarily reflect client views leaves open arguments over framing that could theoretically warrant a pause.

Yet the broader context makes a purely procedural explanation feel incomplete. The report's core prescriptions are politically sensitive: rejecting "absolute versions of victory," explicitly accepting the legitimacy of the Chinese Communist Party, urging mutual recognition of nuclear deterrence, and suggesting that the United States should, in practice, discourage some Taiwanese actions that risk provoking Beijing. These themes run counter to dominant public narratives that frame U.S. policy as a moral struggle against a hostile authoritarian regime, with Taiwan presented primarily as a frontline democracy to be unconditionally defended.

Seen in that light, the withdrawal looks less like the correction of a technical flaw and more like a symptom of political discomfort. Keeping the PDF accessible preserves the integrity of the research record for those who go looking. Removing the full report from RAND's searchable catalog and replacing it with a minimalist "withdrawn" notice reduces its visibility at precisely the moment when commentators across the political spectrum are beginning to discuss it. It is a half-step: not outright censorship, but a signal that the institution is uneasy about being too publicly associated with a strategy of managed coexistence.

In historical perspective, this is striking. Early megadeath-era reports, with their charts of millions dead, could circulate under RAND's banner with comparatively little self-consciousness. A document that calls for accepting limits, codifying mutual deterrence, and avoiding catastrophic miscalculation in the U.S.–China rivalry apparently cannot. That asymmetry is part of what this paper seeks to understand.

4. From World Targets in Megadeaths to Interdependent Vulnerabilities

The old nuclear maps treated the world as a set of discrete, physical targets: cities, bomber bases, missile silos, ports, industrial complexes. Each dot on the map could, in principle, be "removed" by a single warhead, and the central question was how many dots could vanish before a state ceased to exist in any meaningful sense. The megadeath mindset translated this geometry into casualty curves and "damage levels."

The contemporary U.S.–China rivalry operates on a very different landscape. The most consequential targets are no longer only cities or bases but networks: fabrication plants that anchor global supply chains, undersea cables that carry the world's data, shipping chokepoints through which energy and trade flow, and digital platforms whose failure would ripple instantly across the planet. These are not merely national assets; they are nodes in systems that sustain everyday life from Los Angeles to Lagos.

The shift from physical concentration to systemic interdependence means that war planning can no longer pretend that damage will be neatly bounded by national borders. A strike on one country's semiconductor facilities radiates through global industry; a disruption of a key shipping lane drives up prices and insecurity far from any battlefield. The "world targets in megadeaths" have been overlaid by "world targets in vulnerabilities," and the latter are, by design, shared.

4.1 Replacing city lists with fabs, cables, ports, and platforms

One way to understand the change is to imagine the targeting binder that would sit on the table of a contemporary war planner. The classic Cold War volume might have lists of metropolitan areas, industrial regions, airfields, and command bunkers. The twenty-first-century equivalent would still contain some of those, but alongside them would be entirely different classes of target.

First are fabrication plants and specialized industrial facilities, especially in the semiconductor sector. These are geographically concentrated, capital-intensive, and extremely difficult to replace quickly. Damage to a small number of such sites would disrupt production of chips that

underpin everything from smartphones to weapons systems. Here, the “target” is not just a building; it is an entire layer of the world economy that depends on that building operating.

Second are communication infrastructures, particularly undersea cables, landing stations, and major data centers. Severing or disabling a few key cables can cut off significant portions of global internet traffic. Disrupting large data centers or cloud providers has cascading effects on finance, logistics, and even basic public administration.

Third are ports and logistical hubs. Container terminals, oil and gas facilities, and railway chokepoints have become tightly integrated into just-in-time supply chains. Damage to a few crucial hubs can ripple along shipping schedules, inventories, and pricing structures worldwide.

Finally, there are digital platforms: payment systems, messaging networks, trading platforms, and other software infrastructures. These are less obviously “targets” in the conventional sense, but in practice they function as critical nodes. Cyber operations that degrade or corrupt them can produce effects comparable to physical attacks, albeit with less visual drama.

The point is not that cities have ceased to matter, but that the strategic picture now includes a dense mesh of shared vulnerabilities. Any realistic assessment of U.S.–China conflict must recognize that many of the most attractive targets are precisely those whose destruction would harm both sides and innumerable third parties.

4.2 Economic entanglement and mutual hostage situations

Interdependence has often been sold as a force for peace. When economies are tightly entangled, the argument goes, the costs of war become prohibitive. The U.S.–China relationship illustrates a more ambiguous reality. Deep economic ties do indeed raise the price of open conflict, but they also create a peculiar form of mutual hostage-taking.

Both economies depend on access to each other’s markets, inputs, and technologies in complex ways. American firms rely on Chinese manufacturing capacity and consumer demand; Chinese firms rely on access to U.S. technology, financial systems, and high-value markets. Global supply chains bind these relationships into wider systems: European, Asian, African, and Latin American economies all sit downstream of U.S.–China trade and investment flows.

In this context, threats to weaponize interdependence become a central part of strategic competition. Sanctions, export controls, investment bans, and selective decoupling are used to signal displeasure, coerce behavior, or slow the rival’s advance. At the extreme, planners contemplate cyber operations or physical disruptions targeting the other side’s economic lifelines.

Yet the same structures that make such strategies attractive also make them dangerous. Because supply chains are global, harm rarely remains confined to the intended target. A

concerted campaign against another state's shipping, manufacturing base, or financial connectivity risks triggering wider crises: shortages, inflation, unemployment, and political instability around the world. The mutual hostage situation thus extends far beyond the two main actors.

The RAND stabilization report's focus on semiconductors, critical infrastructure, and technological decoupling must be read against this backdrop. It implicitly acknowledges that the old notion of "acceptable damage" makes little sense in a deeply entangled system. There is no obvious way to punish an adversary's economy at scale without also inflicting serious collateral damage on oneself and on third parties.

4.3 Information, finance, and supply chains as non-kinetic escalation ladders

In the megadeath era, escalation ladders were mostly defined in terms of military actions: alerts, mobilizations, conventional strikes, limited nuclear use, and finally large-scale exchanges. Today, information flows, financial instruments, and supply chain disruptions form parallel ladders that can be climbed without firing a shot.

Information operations range from propaganda and disinformation campaigns to cyber intrusions that steal or leak sensitive data. They can be used to embarrass leaders, sow distrust among allies, or undermine domestic cohesion. Escalation along this ladder involves more aggressive attacks on critical information infrastructure, such as media networks, electoral systems, or emergency services.

Financial measures constitute another ladder. States can impose targeted sanctions on individuals, broaden them to cover sectors, restrict access to capital markets, freeze assets, or cut institutions off from key payment systems. Each step raises the cost to the targeted state and invites retaliation, whether through countersanctions, diversification away from the sanctioning state's currency, or efforts to build parallel financial architectures.

Supply chain disruptions form a third non-kinetic ladder. Initial steps might involve export controls on specific technologies, administrative delays, or tighter regulatory scrutiny. Higher rungs could include embargoes, seizures of cargo, or deliberate interference with shipping routes and logistics. The top of this ladder would be outright blockades or covert attacks designed to cripple key flows of goods.

These ladders interact with each other and with military postures. A cyberattack on financial infrastructure might be interpreted as a strategic strike; a shipping disruption in a contested sea lane could trigger naval brinkmanship. Because the effects of non-kinetic actions are diffuse and sometimes delayed, misperceptions and unintended escalation are common risks. Stabilization in this environment requires not only military-to-military communication channels but also

shared understandings of how far escalation in information, finance, and supply chains can go before it is treated as an act of war.

4.4 Why contemporary war planning cannot be separated from global systems survival

The cumulative effect of these changes is that contemporary war planning cannot be cleanly separated from questions of global systems survival. In the megadeath framework, strategists could, at least in theory, imagine a limited nuclear exchange that left "distinguishable postwar states" with surviving populations and infrastructure. The moral adequacy of such thinking is deeply questionable, but its internal logic presupposed that the damage, however horrific, would be bounded in space and time.

Under conditions of dense interdependence, this assumption breaks down. A major conflict between the United States and China, even if confined to conventional weapons and non-kinetic measures, would likely produce cascading failures across financial, logistical, and technological systems worldwide. The resulting crises—resource shortages, economic collapses, political breakdown—could in turn feed back into renewed conflict and further breakdown. There is no obvious postwar "steady state" to point to, only scenarios of long-term disruption.

This does not mean that great powers will automatically avoid conflict. It means that any serious planning must grapple with the possibility that the line between "war over there" and "system collapse everywhere" has grown perilously thin. Military options cannot be evaluated solely in terms of battlefield outcomes or bilateral balances of power. They must be assessed in terms of their impact on the shared infrastructures that make contemporary life possible.

The withdrawn RAND report, by emphasizing mutual restraint around critical technologies and infrastructure, gestures toward this reality. It implicitly recognizes that some forms of "victory" in a U.S.–China conflict would be indistinguishable from defeat at the level of global systems. To stabilize the rivalry is thus not only to reduce the risk of conventional or nuclear war but to acknowledge that the continued functioning of the world's networks, platforms, and supply chains is itself a vital interest—one that cannot be neatly assigned to one flag or another.

5. Strategic Culture and Political Red Lines in Washington

The fate of the withdrawn stabilization report cannot be understood purely as a technical episode. It sits inside a broader strategic culture in Washington that still prefers narratives of eventual victory to permanent coexistence, that clothes power politics in moral exceptionalism, and that is tightly wired into domestic electoral incentives. In this environment, serious proposals for managing rivalry with China, rather than "winning" it, easily acquire the smell of appeasement.

The contrast with the megadeath era is instructive. Then, brutal technical candor about nuclear devastation could coexist with public rhetoric of strength because the technocratic core of policy was relatively insulated from mass politics. Now, the boundary between inside and outside is thin. Think-tank reports circulate instantly on social media, get weaponized in partisan struggle, and become litmus tests for being "soft" or "tough" on China. The same report that might have functioned quietly as an internal guide in an earlier era must now survive a culture that treats compromise language as betrayal.

5.1 Victory narratives, moral exceptionalism, and domestic electoral pressures

American grand strategy has long been shaped by victory narratives. From the Second World War through the Cold War, political leaders cultivated a story in which the United States ultimately prevails over its adversaries because it stands for freedom, democracy, and human rights. Defeats or stalemates could be reinterpreted as temporary setbacks; the arc of history, in the national imagination, bends toward American success.

This narrative is reinforced by a strong current of moral exceptionalism. The United States is not simply one great power among others. It sees itself, and sells itself domestically, as uniquely virtuous. When rivalries are framed as contests between good and evil, "managing" an adversary's power begins to look like accepting evil as permanent. For many political actors and voters, that is a difficult pill to swallow.

Electoral pressures magnify this difficulty. Few politicians want to be accused of weakness toward a rising authoritarian superpower. Campaign consultants test messages that promise toughness on China; interest groups and media ecosystems reward alarmist framing. Proposals that emphasize limits, shared vulnerabilities, or coexistence do not poll well. Even raising them risks primary challenges, negative advertising, and hostile coverage.

In such a context, a report that calmly accepts the long-term legitimacy of the Chinese state, rejects "absolute victory" as an objective, and focuses on damage limitation in specific domains runs against the grain. It may be intellectually responsible, but it is politically radioactive. The safer route is to keep such thinking behind closed doors and let public documents lean toward more muscular language.

5.2 Hawks, doves, and the suspicion of "legitimizing" the Chinese Communist Party

Within Washington's foreign-policy community, there remains a spectrum of views about China. There are hardline hawks who see the rivalry as an existential ideological struggle and advocate for aggressive containment or even regime change. There are cautious doves, or at least realists, who see China as a durable great power whose system is unlikely to collapse quickly and who therefore prioritize stability over maximalist pressure. Between these poles sit a range of views that mix moral condemnation with pragmatic constraint.

The stabilization report clearly leans toward the realist end of this spectrum. It treats the Chinese Communist Party as the current governing authority of a major state whose domestic future is not under U.S. control. It emphasizes mutual recognition, mutual deterrence, and negotiated boundaries of competition. It suggests that attempts to deny Beijing any sense of legitimacy are both futile and destabilizing.

For hawks, this kind of posture is almost indistinguishable from conceding the argument. Accepting CCP legitimacy, even as an analytical starting point, is read as normalizing an oppressive regime. Exploring incentives for peaceful approaches to Taiwan's future is read as paving the way for unification under Beijing's terms. Proposing mutual commitments not to target each other's critical infrastructure is read as tying America's hands in a future conflict.

Suspicion of "legitimizing" the adversary is not new; similar accusations were leveled during détente with the Soviet Union. What is new is the speed and intensity with which such suspicion can surge through online spaces and partisan media. A paragraph in a report can be pulled out, circulated without context, and turned into a symbol of cowardice or betrayal. Under these conditions, even modest, carefully worded recommendations for coexistence can trigger disproportionate hostility. The incentive for institutions like RAND is to tread lightly, or retreat.

5.3 Taiwan as symbol, ally, and tripwire for credibility politics

Nowhere are these dynamics more concentrated than in the politics of Taiwan. In strategic terms, Taiwan is a crucial node: a self-governing democracy, a key link in semiconductor supply chains, and a physical location astride vital sea lines. In symbolic terms, it has become an emblem of resistance to authoritarian pressure and a test case for whether the United States stands by its friends.

For many in Washington, Taiwan is not just another security commitment. It is the place where backing down would allegedly shatter American credibility worldwide. If the United States is perceived to abandon Taiwan, the argument goes, allies in Japan, South Korea, and beyond will doubt U.S. guarantees. Adversaries will see an invitation to probe elsewhere. The entire architecture of alliances built after the Second World War is thought to hang, at least rhetorically, on what happens in the Taiwan Strait.

This creates powerful political incentives to talk about Taiwan in absolute terms. Members of Congress travel to Taipei and make high-profile statements of support. Op-ed pages frame the issue as a moral line in the sand. Analysts warn of catastrophic consequences if Washington fails to respond decisively to any move by Beijing. In this environment, nuance is rare.

The stabilization report's discussion of Taiwan cuts against this absolutist mood. By emphasizing deterrence without endorsing formal independence, by suggesting incentives for Beijing to pursue gradual approaches, and by hinting that Washington may sometimes need to discourage

Taipei from symbolic provocations, it reframes Taiwan as part of a three-sided stability problem rather than a one-dimensional litmus test of American resolve. That is analytically sound but ideologically explosive. To the guardians of credibility politics, any suggestion that Taiwan is a variable to be carefully managed, rather than a sacred cause to be defended at any cost, looks dangerously like the first step toward "selling out" an ally.

5.4 Why candid stabilization language is harder to publish than megadeath tables ever were

It may seem paradoxical that frank discussions of millions of deaths were easier to publish than a candid blueprint for stabilizing rivalry. The paradox dissolves when we notice how the surrounding political and media ecosystems have changed.

In the early nuclear era, elite strategic discourse was relatively insulated. A small number of officials, analysts, and journalists mediated the knowledge of war plans. Reports circulated among specialists; the general public caught only glimpses. When technical studies did become public, they did so in a world with fewer, slower channels for outrage. Critics could denounce the moral coldness of megadeath planning, but their denunciations did not immediately turn into primary challenges, fundraising emails, or targeted social-media campaigns.

Today, the boundary between technical and political speech is porous. Any document uploaded to a think tank's website exists, effectively, on the same stage as campaign slogans and partisan attacks. Lines written for an audience of specialists are instantly available to activists, commentators, and foreign propaganda outlets. There is little room for an "inside voice" that differs significantly from the "outside voice."

Under these conditions, straightforward stabilization language is more dangerous than abstract tables about potential casualties. A chart of hypothetical megadeaths can be dismissed as technical modeling, however grotesque. A paragraph stating that the United States should accept the indefinite survival of a rival authoritarian regime, avoid trying to win outright, and exercise restraint over an ally's behavior cannot. It touches raw nerves about identity, morality, and leadership.

Moreover, the contemporary policy marketplace is saturated with actors whose funding and influence depend on sustaining a sense of looming threat. Defense contractors, certain corners of the think-tank world, and segments of the media all benefit from narratives of intensifying confrontation. A report that hints at ceilings on confrontation, that questions the wisdom of aiming for decisive victory, or that recommends mutual restraint on technology warfare runs against these material incentives as well as ideological ones.

The result is a strategic culture in which serious thinking about coexistence exists, but often in shadow form. It appears in war games, in confidential briefings, in classified annexes. It rarely appears in public documents designed to be read by anyone with a browser. The withdrawal of

the RAND stabilization report is thus not an isolated act but a small, telling confirmation of a broader pattern: it is still easier in Washington to describe how the world could burn than to describe, in plain language, how a rival might be lived with.

6. A Logostic Reading: q_US(t), q_PRC(t), and τ_not_war

A logostic reading treats grand strategy not as a static doctrine but as a trajectory. At any given moment, the United States occupies a state q_US(t) in a high-dimensional space of beliefs, institutions, interests, and fears. The People's Republic of China occupies its own state q_PRC(t). Both move through time under internal and external forces. Above and around these contingent states lies a family of possible τ's: truth-manifolds that encode what would actually sustain life, justice, and long-term viability.

In this frame, a world haunted by megadeath planning corresponds to one attractor in the space of q_US(t): a basin in which deterrence by terror, victory narratives, and acceptance of mass destruction as a planning variable all reinforce each other. The emerging language of managed rivalry and interdependence hints at a different attractor: one in which the limits of power, the fragility of shared systems, and the inadmissibility of certain forms of catastrophe are more fully acknowledged. τ_not_war names a particular manifold in this space: the set of strategic postures under which large-scale war between nuclear, systemically central powers does not occur, not by luck but by disciplined constraint.

The withdrawn RAND report can be read as a brief excursion of q_US(t) toward τ_not_war. Its disappearance from public view is then a sign that gravitational pull from the older megadeath attractor remains strong.

6.1 Megadeath doctrine as one historical attractor for q_US(t)

Megadeath doctrine did more than supply tools; it carved channels in the landscape of strategic thought. Once planners accept that millions of deaths can be plotted, compared, and optimized, q_US(t) is nudged into a basin where certain assumptions feel natural. It becomes normal to speak of credibility as a willingness to inflict large-scale destruction, to treat existential risks as levers, and to think of ethics as lying in the careful calibration of horror rather than in its categorical rejection.

In logostic terms, this is an attractor because it is self-reinforcing. The more an institution reasons in megadeaths, the more its models, war games, and career incentives reward those who can navigate that calculus. The more it signals to adversaries that it is capable of rational catastrophe, the more it must continue to perform that role to sustain deterrence. Over time,

q_US(t) settles into a groove where planning for catastrophe becomes synonymous with realism.

This attractor does not require anyone to desire destruction. It only requires them to accept that the path to safety runs through a willingness to threaten it. That acceptance shapes procurement decisions, alliance structures, and rhetorical habits. Even when nuclear planning becomes less prominent, the underlying pattern persists: crises are framed in terms of escalatory ladders, thresholds, and reputational stakes, all anchored in the idea that the ultimate guarantor of order is the credible possibility of extreme violence.

From a logostic standpoint, this megadeath attractor is misaligned with any τ that treats human life and planetary systems as non-fungible. Yet alignment is not binary. q_US(t) can be partially aligned along some coordinates (for example, refusing first use in practice, even if not in doctrine) while remaining curved toward destructive basins along others. The early RAND world normalized that curvature.

6.2 Managed rivalry and interdependence as a competing attractor

The contemporary language of managed rivalry and interdependence represents a different basin of attraction. In this basin, planners recognize that major powers are embedded in shared networks whose destruction would backfire. They see that certain forms of "victory" would destabilize the very systems their own societies depend on. They begin, however tentatively, to treat systemic survival as a constraint, not merely a variable.

For q_US(t), moving toward this attractor involves several shifts. First, the rival is no longer modeled primarily as an object to be deterred or defeated but as a structurally enduring presence: a node that cannot be eliminated without breaking the network itself. Second, escalation ladders are no longer only military; they include cyber, financial, and supply-chain steps whose effects are difficult to bound. Third, the notion of acceptable damage shrinks, because planners admit that complex systems do not recover cleanly from large shocks.

Interdependence plays an ambivalent role. It is both a stabilizer and a source of leverage. On the one hand, mutual dependence discourages outright war by raising costs. On the other hand, it tempts each side to experiment with weaponized disruption. In logostic terms, interdependence deepens the curvature of the landscape: local moves in one domain tilt the manifold in distant places.

The RAND stabilization report, in treating Taiwan, the South China Sea, and critical technologies as domains where mutual restraint is preferable to maximalist pressure, gestures toward this managed-rivalry attractor. It assumes that q_US(t) and q_PRC(t) can, in principle, inhabit a configuration where competition is intense but bounded. It assumes there is a basin in which both states accept that they are co-tenants of vulnerable systems, not sovereigns over infinitely

resilient terrain. τ_not_war is the idealized boundary of that basin: the set of trajectories in which rivalry never crosses into system-damaging war.

6.3 Curvature, humility, and the acceptance of finite power under τ_not_war

Logostics is concerned with curvature: the way in which local decisions bend trajectories toward or away from a given τ. Under τ_not_war, curvature points away from brinkmanship and toward restraint. It rewards humility: an honest admission of finite knowledge, finite control, and finite legitimacy. For a state habituated to megadeath logic and moral exceptionalism, this is a difficult posture to adopt.

Humility in this context does not mean passivity. It means accepting that even a superpower cannot dictate outcomes unilaterally without destabilizing the system it inhabits. It means acknowledging that allies have agency, adversaries have their own narratives of security and humiliation, and complex infrastructures are not easily revived after being struck. It also means accepting that some goals are unattainable at tolerable cost: permanent dominance, universal ideological homogeneity, or the elimination of all threats.

Under τ_not_war, a healthy q_US(t) would internalize these limits. It would frame deterrence as a means of buying time and space for coexistence, not as a ladder to eventual victory. It would treat arms control, crisis hotlines, and mutual non-targeting agreements not as concessions but as structural necessities. It would also recognize the moral asymmetry between threatening megadeath and refusing to do so, even at some cost in bargaining leverage.

Curvature toward this τ requires institutional changes. Promotion systems must reward caution and prudence, not only boldness. Public narratives must make room for stories of managed stalemate, not only iconic triumphs. Legal and ethical frameworks must treat certain actions as unthinkable, not merely costly. In practice, q_US(t) is pulled in two directions: legacy incentives tug toward displays of dominance, while an emerging awareness of interdependence and fragility tugs toward humility. The stabilization report briefly aligns with the latter pull.

6.4 How the withdrawal reflects friction between institutional q and any stabilizing τ

The withdrawal of the RAND report is a small but vivid instance of friction between an institutional q and a stabilizing τ. For a moment, RAND's output aligned more closely with τ_not_war: it explicitly rejected absolute victory, emphasized mutual recognition, and centered systemic vulnerabilities in its recommendations. In doing so, it nudged q_US(t) toward a basin where coexistence, not dominance, is the primary objective.

The reaction suggests that this alignment collided with other curvatures in the system. Political narratives that demand toughness, organizational cultures invested in primacy, and economic interests tied to a high-threat environment all exert force on q_US(t). They make certain

sentences feel dangerous: accepting the legitimacy of the rival, acknowledging ceilings on what force can achieve, or proposing limits on technological warfare. When those sentences appear in a public document, feedback loops activate. Commentators amplify them, adversarial media weaponizes them, and sponsors worry about association.

From a logostic perspective, the "further review" label is not neutral. It is a small act of vector correction: a way of reducing the visible component of movement toward τ_not_war without necessarily invalidating the analysis. The underlying ideas may continue to influence internal discussions, classified planning, or quiet dialogues. But the public trajectory is pulled back toward a safer, more ambiguous position, one that does not so clearly contradict the dominant story of eventual victory over an unredeemable adversary.

This friction reveals the gap between what institutions dimly know and what they are prepared to say. At some level, many actors understand that a war between the United States and China, fought across physical and technological domains, would be catastrophic for the global system. They understand that τ_not_war is not a luxury but a condition for continued life as they know it. Yet q_US(t), as embodied in concrete organizations, remains entangled with habits, interests, and myths that resist full alignment.

The withdrawn report is thus a small parable of misalignment. It shows an institution briefly articulating a more honest relationship to limits and then recoiling from its own candor. In logostic terms, the episode reminds us that aligning q(t) with any life-preserving τ is not a one-time declaration but an ongoing struggle against the inertia of older attractors.

7. After Megadeaths: Ethics, Policy, and the Possibility of Repentance

To talk about strategy "after megadeaths" is to admit that something in the earlier mindset was not merely dangerous but morally deformed. It is not enough to update models, rename doctrines, or shift priorities from nuclear salvos to supply-chain resilience. If great powers have spent decades treating the destruction of cities as a bargaining chip, and now treat systemic economic collapse or technological breakdown as a tolerable risk in pursuit of advantage, then a deeper question presses in: what would repentance look like at the level of states and strategic cultures?

Repentance, in this context, does not mean a sentimental declaration of regret. It means a turning: a shift in what counts as admissible, in how leaders speak about risk, and in what they are prepared to do or refuse to do. It means naming certain practices as incompatible with any decent understanding of human dignity, even if those practices promise leverage. It also means constructing structures that make backsliding harder: laws, norms, alliances, and habits of discourse that support restraint rather than glorify brinkmanship.

The withdrawn stabilization report, with its emphasis on coexistence, mutual restraint, and the protection of shared infrastructures, hints at what a post-megadeath ethic might require. But the report's retreat from public view shows how contested such an ethic remains. This section sketches minimal conditions for honest strategy today, explores what a non-victory paradigm would demand for alliances and Taiwan, considers guardrails for science and technology, and asks what roles citizens, churches, and dissident networks can play in resisting escalation culture.

7.1 Minimal ethical conditions for honest great-power strategy today

The first ethical condition is intellectual honesty about consequences. Any serious great-power strategy must confront, in plain language, what large-scale war would do in a world of nuclear arsenals, dense interdependence, and fragile infrastructures. This includes not only direct casualties but also long-term effects on food systems, public health, climate, financial stability, and governance. Strategies that treat such risks as abstract "costs" to be balanced against "gains" fail at the threshold. An honest doctrine starts by saying: there are forms of conflict whose likely consequences are indefensible, regardless of who "wins."

A second condition is the rejection of deliberate mass targeting of civilians as a legitimate instrument of policy. This is easy to affirm in principle and hard to practice in an era of dual-use infrastructure and integrated systems. The temptation will always be to argue that crippling an adversary's economy or networks is "mere coercion." A minimal ethic insists that when planners know an action will predictably throw millions into poverty, hunger, or chaos, they cannot wash their hands with euphemisms. Civilian harm on that scale must be treated as morally central, not incidental.

Third, there must be acknowledgment of limits. No state, however powerful, can guarantee its own perfect security, eliminate all threats, or preserve its primacy indefinitely without courting disaster. An ethical strategy accepts that some ambitions are structurally incompatible with global survival. Pursuing unconditional dominance, insisting on unilateral freedom of action everywhere, or framing every concession as surrender are not just imprudent; they are incompatible with treating other peoples as anything more than obstacles.

Fourth, there must be procedural transparency toward one's own citizens. While operational details can remain classified, the broad contours of risk cannot. Populations subjected to the dangers of nuclear deterrence, technological escalation, and economic warfare have a right to understand the stakes. This requires public communication that does not hide behind clichés but explains why certain paths are closed off and why others, however unsatisfying, are chosen.

Finally, there must be some commitment to reciprocity. If one state claims that certain classes of attack are unacceptable, it must be willing to renounce them itself. Ethical strategy cannot

consist merely in demanding restraint from others while reserving every option for oneself. Minimal reciprocity is not naive symmetry; it is a recognition that norms have no force if they are transparently one-sided.

7.2 Rethinking alliances, extended deterrence, and Taiwan under a non-victory paradigm

A non-victory paradigm starts from the premise that the central aim of U.S.–China strategy is not to achieve decisive dominance but to avoid systemic catastrophe while preserving as much space as possible for free societies to survive and flourish. This does not abolish alliances or extended deterrence, but it reframes them.

Alliances, under such a paradigm, are no longer promises to prevail in every conceivable confrontation. They are commitments to mutual defense within the limits of what can be done without triggering world-scale breakdown. That means allies must be spoken to as adults, not as props in global narratives. They must be told, candidly, where the hard ceilings lie: which scenarios would almost certainly escalate beyond control, which symbolic acts are not worth the risk, and what kinds of burden-sharing and restraint are necessary to keep security guarantees credible rather than suicidal.

Extended deterrence must shed some of its absolutist rhetoric. The point of deterrence is not to convey a willingness to burn down the world. It is to shape expectations so that no actor believes it can gain by crossing certain thresholds. In practice, this may mean narrowing the range of contingencies for which automatic military responses are promised, while expanding support in non-military domains. Assisting allies in building resilient economies, robust civil defenses, and independent diplomatic channels can reduce their perceived need to test the outer limits of security guarantees.

Taiwan is the hardest case. In a non-victory paradigm, the question is not “How can the United States ensure that Taiwan’s political status is never altered by pressure from Beijing?” but “How can the United States help ensure that twenty million people on a heavily armed island do not become the center of a war that wrecks the global system?” That shift does not abandon Taiwan; it changes what “standing with Taiwan” means.

It may mean prioritizing the island’s survival and self-governance over maximalist symbolic goals. It may mean strengthening Taiwan’s defensive capabilities while quietly discouraging gestures that force Beijing into a corner. It may mean exploring arrangements that preserve de facto autonomy and democratic life without insisting on formal declarations that Beijing has vowed to treat as casus belli. It may also mean helping Taiwanese society prepare for a long-term reality in which neither unification nor independence is cleanly achievable, but in which a messy, sometimes humiliating status quo is, tragically, the least harmful option.

For Washington, a non-victory paradigm demands a new honesty with its own public. Politicians must stop implying that every ally can be defended at any cost without grave risk, and that any compromise over symbolic issues equals betrayal. Citizens must be told that there are scenarios in which pressing a point of principle to the brink would imperil not only one island but the entire mesh of systems that sustain life around the world. In such scenarios, restraint is not cowardice; it is sanity.

7.3 Guardrails for science and technology competition that stop short of megadeath by other means

Science and technology competition between the United States and China is not going away. The question is whether it proceeds with any guardrails, or whether it becomes the main avenue through which megadeath logic returns in new forms: pandemics sparked by reckless research, cascading blackouts from cyberattacks, financial collapse triggered by deliberate algorithmic sabotage.

Guardrails begin with clear “no-go” zones. States should identify classes of technological activity that are simply too dangerous to be turned into weapons in a rivalry: deliberate erosion of each other’s public health infrastructure, attacks on nuclear safety systems, systemic sabotage of electrical grids, and unbounded development of autonomous lethal systems are examples. Formal treaties may be hard to negotiate, but unilateral declaratory policies, mirrored by the rival and reinforced by track-two dialogues, can begin to crystallize norms.

Next, there must be transparency mechanisms. Joint technical working groups, even in tense times, can help clarify what is happening in high-risk areas like AI, synthetic biology, and cyber operations. The point is not to hand over secrets but to lower the risk of misinterpretation. If both sides understand the difference between a routine software update and a preparations-for-attack pattern, they are less likely to misread noise as escalation.

Third, there should be crisis communication channels specific to technology. A catastrophic, ambiguous incident—a massive outage, a sudden financial disruption, an AI system behaving unpredictably—could be read as an attack when it is not, or as the opening salvo of a broader campaign when it is meant as a one-off signal. Dedicated hotlines, pre-agreed notification protocols, and standing expert teams who can consult across borders in real time are essential.

Fourth, guardrails require some insulation of the global commons. Not every element of the technological stack can be protected, but certain layers are obvious candidates: basic pandemic surveillance networks, limited financial backbones needed for humanitarian trade, and agreed safe modes for critical communication systems. Even adversaries have an interest in preserving these during crises, because their destruction would make coordinated de-escalation nearly impossible.

Finally, guardrails demand domestic discipline. States must resist the temptation to push up to the edge of every possible advantage. This may mean declining to pursue certain capabilities, even if technically feasible, because their existence would make panic or pre-emption more likely in a crisis. It may also mean subjecting strategic technological programs to independent ethical review, including voices that are not financially or institutionally invested in escalation.

Without such guardrails, science and technology competition risks becoming megadeath planning by other means: the cultivation of systemic vulnerabilities that, when triggered, produce human suffering and disorder on a vast scale without the visual clarity of a mushroom cloud.

7.4 Roles for citizens, churches, and dissident networks in resisting escalation culture

The culture that makes megadeath thinking, or its modern equivalents, seem normal does not arise only from think tanks and war colleges. It is sustained by stories: about enemies, about honor, about what counts as courage. Resisting escalation culture therefore cannot be left to experts alone. Citizens, religious communities, and dissident networks have their own essential roles.

Citizens can insist on truthfulness. They can demand that leaders explain the real risks of confrontation rather than hiding behind slogans. They can support journalists and scholars who map the systemic consequences of war, not just its tactical dimensions. They can refuse to reward politicians who invoke foreign threats as all-purpose justification for domestic failures. In democracies, public appetite for nuance and restraint can, over time, alter what strategists feel permitted to say.

Churches and other religious communities can reframe moral imagination. They can remind congregants that people on the other side of a border are not abstractions but bearers of the same image, the same dignity. They can teach that there are lines of cruelty and recklessness that cannot be crossed in pursuit of national advantage, no matter how loudly the surrounding culture clamors for toughness. They can host conversations that connect scriptural or philosophical teachings on peace and justice with concrete issues like Taiwan, sanctions, and technological warfare, rather than treating geopolitics as a remote, secular domain.

Dissident networks—loosely organized communities of thinkers, activists, veterans, technologists, and believers—can provide alternative channels of perception. Across national boundaries, they can share information that state media suppress, debunk demonizing narratives, and cultivate empathy with ordinary people in rival states. Within their own societies, they can challenge the habits of speech that make genocidal or system-crushing options seem thinkable. They can also experiment with new forms of dialogue, using digital tools to connect those who refuse to see their counterparts as inevitable enemies.

Repentance, at this level, becomes communal. It looks like activists admitting when their own side has escalated unnecessarily. It looks like veterans speaking honestly about the costs of past wars and refusing to romanticize them. It looks like religious leaders confessing their own complicity when they have blessed campaigns that could only end in ruin. It looks like citizens acknowledging their susceptibility to flattery and fear, and choosing, however imperfectly, to align their voices with those who are trying to bend the trajectory toward non-catastrophic outcomes.

In logostic terms, each of these actors carries a small piece of q(t): the time-dependent state of a community's beliefs and habits. Their choices can subtly shift curvature toward or away from τ_not_war. No individual action can force alignment. But a critical mass of small refusals—refusals to celebrate escalation, to dehumanize rivals, to treat systemic destruction as an acceptable bargaining move—can change what is sayable, and thus what is thinkable, in the halls where grand strategy is made.

After megadeaths, the central question is no longer how cleverly a state can navigate escalation ladders. It is whether societies can build enough moral ballast to say, and to mean, that certain ladders will not be climbed at all. Ethics, policy, and repentance are not three separate tasks. They are facets of the same work: teaching powerful institutions, and the cultures that sustain them, to live within limits that preserve the possibility of a shared future.

8. Conclusion – What RAND's Two Eras Reveal

The juxtaposition of RAND's nuclear-age megadeath planning and its contemporary attempt to stabilize the U.S.–China rivalry reveals more than an institutional history. It exposes a deep tension in how a superpower imagines its own role in a fragile world. In one era, RAND helped teach Washington to think in units of millions dead, wrapping that horrific capacity in a language of rational deterrence and escalation control. In another, RAND briefly articulated a strategy that accepts the permanence of a rival, foregrounds mutual vulnerability, and seeks to avoid catastrophic war in concrete domains such as Taiwan, the South China Sea, and critical technologies.

Seen together, these two RANDs tell a story about the shifting object of strategy. The early object was victory, even if that meant surviving a nuclear exchange in a "distinguishable postwar state." The newer object, at least in the withdrawn report, is continued coexistence under conditions of deep interdependence, where victory in the old sense would amount to systemic self-harm. Yet the uneasy treatment of this newer posture—released, noticed, then withdrawn for "further review"—shows that the older gravity field of megadeath thinking and victory narratives has not lost its pull.

A logostic reading frames these as different attractors for q_US(t), the time-dependent state of American strategic culture, and for q_PRC(t), its Chinese counterpart. τ_not_war represents a manifold of strategic truths in which large-scale war between system-defining powers does not occur, because key actors have internalized the catastrophic consequences and adjusted their behavior. RAND's two eras are then glimpses of how close, and how far, a powerful state can move toward that manifold.

8.1 Summary of the argument and its limits

This paper has traced a line from "World Targets in Megadeaths" to the withdrawn "Stabilizing the U.S.–China Rivalry" report. It began by sketching how early RAND, embodied in figures like Herman Kahn, built an intellectual framework where megadeaths were acceptable units of planning. Escalation ladders, "tragic but distinguishable postwar states," and the cold rationality satirized in Dr. Strangelove all emerged from a conviction that only by thinking the unthinkable could catastrophe be deterred.

The paper then turned to the contemporary report that seeks not absolute victory but bounded rivalry. It highlighted how this document treats Taiwan, the South China Sea, and science and technology as domains where mutual restraint and crisis management are more realistic objectives than decisive triumph. It noted how the report foregrounds interdependent vulnerabilities—fabs, cables, ports, platforms—and implicitly recognizes that the destruction of these nodes would harm not only the adversary but the global system as a whole.

Next, the analysis examined Washington's strategic culture: victory narratives, moral exceptionalism, domestic electoral pressures, and the suspicion that accepting any legitimacy for the Chinese state amounts to appeasement. In this context, the report's rejection of "absolute victory," its acceptance of long-term coexistence, and its hints of quiet restraint on Taiwan are politically explosive. The decision to withdraw the report's public availability was read as a small but telling concession to that culture.

A logostic perspective was then applied: megadeath doctrine as one historical attractor for q_US(t), managed rivalry and interdependence as another, τ_not_war as the manifold on which catastrophic war is actually avoided. From that standpoint, the stabilization report looks like a brief alignment toward τ_not_war, and its withdrawal looks like friction from older attractors pulling q_US(t) back.

Finally, the paper outlined minimal ethical conditions for honest great-power strategy, rethought alliances and extended deterrence under a non-victory paradigm, proposed guardrails for science and technology competition, and suggested roles for citizens, churches, and dissident networks in resisting escalation culture.

The argument has limits. It necessarily simplifies the internal diversity of both American and Chinese strategic thinking. It treats RAND as a symbol of U.S. elite reasoning, even though many other institutions and actors shape policy. It leans heavily on a single report whose internal drafting dynamics and sponsor interactions remain opaque. It also assumes that τ_not_war can be meaningfully approximated in practice, when in fact human institutions are messy, divided, and prone to self-deception. This is not a blueprint; it is a conceptual sketch meant to clarify stakes and patterns.

8.2 Open questions for strategists, ethicists, and theologians

For strategists, the central open question is whether they can integrate interdependence and systemic fragility into their core models, rather than treating them as afterthoughts. What does deterrence look like when critical infrastructures are shared and when economic warfare can reverberate globally? How should escalation ladders be redesigned to capture not only military rungs but financial, cyber, and supply-chain rungs whose interactions are poorly understood? Can alliance commitments be framed in ways that are both credible to partners and honest about systemic ceilings?

Another strategic question concerns institutional design. What structures inside defense ministries, foreign offices, and intelligence agencies help a state move toward τ_not_war, rather than backsliding into megadeath-style thinking dressed in new clothes? Are there ways to embed “red teams” for systemic risk inside planning processes, so that every proposed strategy is stress-tested against global stability, not just national advantage?

For ethicists, the questions cut deeper. Is it ever morally acceptable to threaten actions that would predictably produce suffering and disorder on a planetary scale, even if one hopes never to carry them out? How should intentions and credible threats be evaluated in a world where signaling resolve often means rehearsing atrocities in planning documents? What does just war thinking become when the distinction between battlefield and home front dissolves into a web of infrastructures?

Ethicists must also grapple with repentance in a collective sense. How can communities move from merely regretting past doctrines to restructuring the incentives that made those doctrines attractive? What forms of public acknowledgment, legal change, and institutional reform are necessary to name megadeath logic as a sin against humanity rather than merely a technical phase?

For theologians, there are yet other layers. How should traditions that speak of powers and principalities interpret strategic cultures that normalize systemic destruction? What does it mean to proclaim peace in an era when peace sometimes looks like accepting humiliating limits on power and deferring cherished aspirations to avoid catastrophe? Can concepts such as sin,

idolatry, and repentance shed light on the worship of victory narratives and the sacrifice of distant lives on the altars of credibility and primacy?

Theologians may also ask how τ_not_war relates to their understanding of providence and judgment. Is a world where superpowers learn to live within hard limits a gift, a discipline, or both? And how should faith communities respond when their own symbols and language are mobilized to sanctify escalatory policies?

8.3 Final reflection: living under τ_not_war in a world still haunted by megadeath thinking

To live under τ_not_war is to accept that great-power war in its full, systemic sense is no longer a morally or practically acceptable tool of policy. It is to treat the avoidance of such war not as a background hope but as a primary constraint on grand strategy. It is to acknowledge that the infrastructures, ecosystems, and social fabrics that sustain billions are too fragile to be subjected to deliberate large-scale shocks without imperiling everyone.

Yet megadeath thinking has not vanished. Its habits linger in metaphors of credibility, in casual talk about "taking out" critical nodes, in the way loss-of-life numbers appear in briefing slides as tidy rows. Its ethos survives in the reflex that equates restraint with weakness and that treats the continued existence of a rival system as an intolerable affront. The ghosts of "World Targets in Megadeaths" still hover at the edges of planning rooms, now wearing the masks of cyber tools, economic sanctions, and technological sabotage.

Living under τ_not_war in such a world will feel, for many, like living in tension. On one side are demands for decisive action, moral clarity, and uncompromising defense of values. On the other are the hard facts of nuclear arsenals, interdependent vulnerabilities, and the likelihood that a major U.S.–China conflict would shred the systems on which everyone depends. The task is not to dissolve this tension with slogans but to inhabit it truthfully.

That may mean accepting that some victories will never be won in the old sense. It may mean telling allies that the world cannot be made safe on their preferred terms without risking far greater harm. It may mean telling one's own citizens that heroism now looks less like charging the hill and more like refusing to start fires that cannot be contained.

The withdrawn RAND report offered, however imperfectly, a glimpse of what it might look like for an influential institution to speak in the language of τ_not_war: acknowledging shared vulnerabilities, rejecting fantasies of total victory, and recommending bounded, disciplined competition. Its retreat from public view is a reminder that such speech is still contested, that the attractor of megadeath logic has not lost its power.

Whether q_US(t), q_PRC(t), and the wider constellation of states can move more fully toward τ_not_war remains an open question. The answer will not depend only on classified documents

and high-level summits. It will depend on what ordinary people are willing to accept in their name, on what kind of greatness they demand from their leaders, and on whether enough voices are prepared to say, clearly and without cynicism, that there are ladders they will not climb, no matter how loudly the old ghosts call.

9. References

Allison, G. (2017). *Destined for War: Can America and China Escape Thucydides's Trap?* Houghton Mifflin Harcourt.

Blackwill, R. D., & Campbell, K. M. (2016). *Xi Jinping on the Global Stage: Chinese Foreign Policy under a Powerful but Exposed Leader.* Council on Foreign Relations.

Freedman, L. (2003). *The Evolution of Nuclear Strategy* (3rd ed.). Palgrave Macmillan.

Herman, A. (2012). *Douglas MacArthur, American Warrior.* Random House.

Homer-Dixon, T. (2000). *The Ingenuity Gap: How Can We Solve the Problems of the Future?* Knopf.

Jervis, R. (1989). *The Meaning of the Nuclear Revolution: Statecraft and the Prospect of Armageddon.* Cornell University Press.

Jervis, R. (1997). *System Effects: Complexity in Political and Social Life.* Princeton University Press.

Kahn, H. (1960). *On Thermonuclear War.* Princeton University Press.

Kahn, H. (1962). *Thinking About the Unthinkable.* Horizon Press.

Keohane, R. O., & Nye, J. S. (2012). *Power and Interdependence* (4th ed.). Longman.

Mearsheimer, J. J. (2001). *The Tragedy of Great Power Politics.* W. W. Norton.

Mazarr, M. J., Kerrigan, A., & Lenain, B. (2025). *Stabilizing the U.S.–China Rivalry.* RAND Corporation.

Niebuhr, R. (1932). *Moral Man and Immoral Society: A Study in Ethics and Politics.* Charles Scribner's Sons.

Perrow, C. (1999). *Normal Accidents: Living with High-Risk Technologies* (Updated ed.). Princeton University Press.

Schelling, T. C. (1960). *The Strategy of Conflict.* Harvard University Press.

Schelling, T. C. (1966). *Arms and Influence.* Yale University Press.

Snyder, G. H. (1961). *Deterrence and Defense: Toward a Theory of National Security.* Princeton University Press.

Waltz, K. N. (1979). *Theory of International Politics.* Addison-Wesley.

Walzer, M. (1977). *Just and Unjust Wars: A Moral Argument with Historical Illustrations.* Basic Books.

Yoder, J. H. (1972). *The Politics of Jesus.* Eerdmans.

Dr. Strangelove or: How I Learned to Stop Worrying and Love the Bomb. (1964). Directed by S. Kubrick. Columbia Pictures.

The author has published approximately 4,600 AI-assisted papers at: https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.